

AD-A156 692

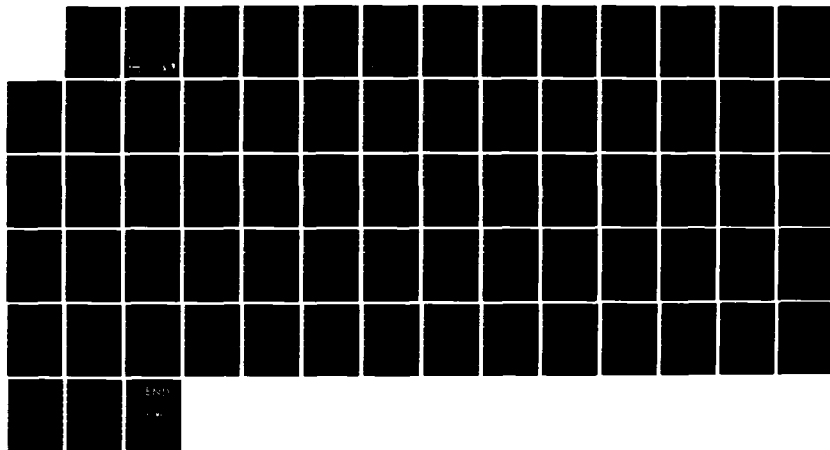
THE DYNAMIC RETENTION MODEL(U) RAND CORP SANTA MONICA
CA R L FERNANDEZ ET AL. APR 85 RAND/N-2141-MIL
MDA903-83-C-0047


1/1

UNCLASSIFIED

F/G 14/1

NL





MICROCOPY RESOLUTION TEST CHART
 NATIONAL BUREAU OF STANDARDS-1963-A

AD-A156 692

FULL COPY

DTIC
ELECTRONIC

15 7 00 333

The National Association of Public Health
Officers of the American Medical Association
Washington, D. C.

The Rand Publications Series: The Report is the product of a study conducted by the
university and the RAND Corporation. The study was conducted by the RAND Corporation
results. The RAND Report reports other sources of information, including the
general distribution. Publications of the RAND Corporation do not necessarily
reflect the opinions or policies of the sponsors of RAND research.

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER N-2141-MIL	2. GOVT ACCESSION NO. D-1176 (6) ✓	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) The Dynamic Retention Model		5. TYPE OF REPORT & PERIOD COVERED Interim
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Richard L. Fernandez, Glenn A. Gotz, Robert M. Bell		8. CONTRACT OR GRANT NUMBER(s) MDA903-83-C-0047
9. PERFORMING ORGANIZATION NAME AND ADDRESS The Rand Corporation 1700 Main Street Santa Monica, CA. 90406-2138		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Assistant Secretary of Defense Manpower, Installations and Logistics Washington, DC. 20301		12. REPORT DATE April 1985
		13. NUMBER OF PAGES 58
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		16. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public Release: Distribution Unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) No Restrictions		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Models Military Personnel Retention (General) Decision Making		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) See reverse side		

This Note is a nontechnical description of the Dynamic Retention Model (DRM), which facilitates analyses of the costs and benefits of proposed changes in the major policies affecting military personnel, including the level and structure of military compensation, the major parameters of the personnel system, and the nondisability retirement system. The DRM predicts a retention rate for a group at any particular point that depends both on future prospects and on history (promotions, pay, bonuses, etc.) for the group members. Because it models individual decisionmaking, the model can simulate the effects of policy changes that have no analogues in past policies. The authors demonstrate the flexibility of the model by comparing it with the Annualized Cost of Leaving (ACOL) model.

A RAND NOTE

THE DYNAMIC RETENTION MODEL

Richard L. Fernandez, Glenn A. Gotz,
Robert M. Bell

April 1985

N-2141-MIL

Prepared for

The Office of the Assistant Secretary of Defense/
Manpower, Installations and Logistics

Rand

1700 MAIN STREET
P.O. BOX 2338
SANTA MONICA, CA 90406-2338


This document has been approved
for public release and sale; its
distribution is unlimited.

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

PREFACE

This Note was produced under the Officer Retention Modeling project, a part of the study program of Rand's Defense Manpower Research Center. The work was sponsored by the Office of the Assistant Secretary of Defense (Manpower, Installations, and Logistics). The project is building dynamic models for predicting the voluntary retention rates of officers in the four active military services.

The Dynamic Retention Model (DRM), developed earlier at Rand by Glenn A. Gotz and John J. McCall, furnished a starting point for this work. The DRM was chosen because it can deal with a wide variety of compensation and personnel policy changes. The Note describes the model in nontechnical terms, without recourse to mathematical expressions, and compares its structure and versatility with those of competing models. It also indicates where the structure and procedures of the DRM may be modified to suit the objectives of the present project.



Accession For	
NTIS CRASI	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

SUMMARY

Congress and the Executive branch continually propose changes in policies affecting military personnel, including the level and structure of military compensation, the major parameters of the personnel system, and the nondisability retirement system. Analyses of the costs and benefits of such proposed changes have been hampered, however, by the lack of retention and inventory projection models capable of dealing with the substantial departures from past practices that are being considered.

Rand is currently building a set of models to predict officer voluntary retention rates for each of the four active military services. We have chosen the Dynamic Retention Model (DRM), developed earlier at Rand by Glenn Gotz and John McCall, as the basic element of this effort. The DRM has the versatility to deal with the wide variety of policy changes mentioned above. This Note describes the model in nontechnical terms, showing that it affords a straightforward, common-sense depiction of the determinants of retention in an uncertain world, even though its formal structure does appear mathematically complex. The Note also shows that the more commonly used Annualized Cost of Leaving (ACOL) model can be viewed as a special case of the DRM. As it has been applied, the ACOL model embodies an important logical inconsistency. It has been properly used in the past to examine specific retention issues, but its inability to deal fully with many policy changes, including some being proposed for the future, limits its long-term usefulness.

The DRM is a two-level model of retention behavior. At one level, it models rational economic behavior under uncertainty. At the second level, it describes the group dynamics, tracing the implications for aggregate retention rates of individual officers' optimal responses to their environment. The model recognizes that individuals differ persistently in their attachment to the military (their "taste" for military service), and that their stay/leave decisions are affected by their tastes as well as by monetary incentives. It further recognizes that unobservable random shocks may affect decisions. Because an

undesirable shock, such as an unwelcome assignment, can be avoided if the individual leaves the military, the model allows the possibility of future shocks to affect decisions today, rather than simply treating them as disturbance terms in a regression.

The model's versatility derives from three features. First, it incorporates a very rich picture of officers' military futures. The model can accommodate changes in promotion opportunity, promotion timing, and high years of tenure within its input parameters. It predicts that if, for example, tenure constraints were relaxed, retention in earlier years of service would increase because the opportunities facing officers within the military would improve. The ACOL model cannot predict this phenomenon because it typically deals only with "average" career paths. Second, the DRM's recognition of persistent differences among officers gives rise to plausible "backward-looking" retention rates. That is, the retention rate for a group of officers at any particular decision point depends upon who is there to be making a decision--what the distribution of tastes is among them--which in turn depends upon the history of compensation and personnel policies preceding the point. The model would predict that a temporary postponement of military pay increases, for example, would reduce retention in the first year, but result in higher retention rates after the postponed increases were restored. The smaller number who remain through the freeze period will have a higher average taste value than their counterparts in a no-freeze baseline case, another phenomenon that the ACOL model cannot predict.

Third, the model reflects the possibility that unforeseeable circumstances may cause an officer to leave even if he has a strong attachment to the military, or to stay beyond his initial commitment even if he has a strong distaste for military life. Thus, it allows retention rates to be influenced not only by policy changes affecting the marginal officer--the one on the borderline between staying and leaving--but also by changes that affect only some infra-marginal group. A good example of this is given in simulations performed by Gotz, in which the policy change was a series of modest bonuses given to all officers in the fifth through eighth and ninth through twelfth years of service, alternatively with and without active duty service obligations

covering the same periods. The imposition of the obligations more than offsets the retention gain in the year of service before the first bonus, demonstrating the value of flexibility to individuals who cannot perfectly predict the future. Flexibility plays no role in the certain world of the ACOL model; because the obligations would not concern the ACOL's marginal officer, who plans to stay to retirement, the ACOL model would predict that the imposition of the obligations would have no retention effect.

A final example illustrates the importance of all three features. Suppose the military retirement system were changed, with benefits reduced for early retirees (20 to 24 years of service) and increased for later retirees. This is similar to a change recently proposed by the Fifth Quadrennial Review of Military Compensation. Officers with less than 12 years of service when the revision is enacted would be allowed to choose between the current and the new systems. The DRM would predict a one-time increase in retention among these officers as a result of the broader opportunities this choice gives them (the ACOL model would not). It would also predict the proportion choosing each system; those choosing the new system would tend to be officers who expect to retire late because they have high taste for the military and/or believe their promotion chances to grades O-5 and O-6 are high (O-4s can only stay to 24 years of service). The ACOL model does not account for persistent differences among individuals, nor does it deal with the possibility that officers may differ in their military career paths. The DRM would predict that in the long run, the new system would tend to discourage low-taste officers--those who if they stayed would retire at or near 20 years of service--from continuing to retirement. In our specific illustration, retention in the years near the first decision point would probably drop, though not by as much as the ACOL model would predict; the ACOL would ignore the improved retention among officers who if they stayed at all would plan to retire late.

None of the examples prove that the ACOL model could not, in the hands of a careful analyst, produce useful and reasonably accurate predictions in many cases. Some of the errors in the model's predictions that can be identified in theory may not be significant in practice. The model clearly is more limited than the DRM, however, in the range of policies with which it can deal fully.

The DRM's versatility does exact a cost: it requires more data than simpler models, and a more complex estimation and projection methodology. Estimating the two distributions--of tastes and of random shocks--requires longitudinal data: successive observations on each officer. The Defense Manpower Data Center (the source of our data) is accustomed to generating longitudinal files, however, so this does not create a problem. The model is specifically designed for maximum-likelihood estimation, and Gotz and McCall have developed a workable program for estimation that we can now refine. Reestimation should not be necessary for a number of years. Finally, forecasting with the DRM is not the simple reversal of the estimation process that it is with most regression models. Thus, part of the development process for the DRM will be to write a program to generate projections, and DMDC will be provided with programs to update the DRM's data inputs. A planned inventory projection model (IPM), to be integrated with the DRM, will make forecasting particularly easy. The IPM will generate promotion rates and other inputs required by the DRM; the DRM in turn will provide retention rates to the IPM. The result will be projections of officer inventories by grade, year of service, or other appropriate divisions. The costs imposed by the DRM's complexity will thus be incurred only during the model's development and estimation, and will not be felt (except in the form of computer time) by the model's user.

CONTENTS

PREFACE	iii
SUMMARY	v
FIGURES	xi
TABLES	xiii

Section

I. INTRODUCTION	1
II. THEORY	3
The Individual Decision	4
Aggregate Assumptions	19
III. APPLICATIONS	23
Up-or-Out Policy	24
Postponement of Active Duty Pay Increases	26
Officer Bonuses with and without Obligations	27
QRMCI Retirement Proposal	30
Summary	34
IV. DATA, ESTIMATION, AND PREDICTION	36
Data	36
Estimation	45
Prediction	49
V. CONCLUSION	53
BIBLIOGRAPHY	57

FIGURES

1. Probability density of the extreme value function
for maxima 20
2. Comparison of annuity equivalents for the QMOC proposal
and the current retirement system 32

PREVIOUS PAGE
IS BLANK

TABLES

1.	Officer Bonus Amounts and Timing	28
2.	Effects of Officer Bonuses on Year-Group Voluntary Retention Rates	29

PREVIOUS PAGE
IS BLANK

I. INTRODUCTION

Proposals for changes in the major policies affecting military personnel, including the level and structure of military compensation, the major parameters of the personnel system, and the nondisability retirement system, are continuously being proposed in Congress and the Executive branch. Analyses of the costs and benefits of such proposed changes have been hampered, however, by the lack of retention and inventory projection models capable of dealing with the substantial departures from past practices that are being considered. A notable exception is the Dynamic Retention Model (DRM) developed at Rand by Glenn Gotz and John McCall.¹ The DRM predicts a retention rate for a group at any particular point that depends both on future prospects and on history (promotions, pay and bonuses, etc.) for the group members. Because it models individual *decisionmaking*, the model can simulate the effects of policy changes that have no analogues in past policies. For policy initiatives that standard models must treat as structural changes, the Dynamic Retention Model can yield plausible predictions of effects.

Rand is currently building a set of models to predict officer voluntary retention rates for each of the four active military services. We have chosen the DRM as the basic element of this effort because of its versatility.² Apparently standing in the way of the DRM's general acceptance as a valuable policy analysis tool has been a perception that it is too complicated and difficult to understand for everyday use. We believe this to be a misperception, deriving perhaps from the fairly complex mathematical structure of the model as it has previously been described. That mathematical complexity is real, but should be a concern only to the researcher who wants to develop an estimated model,

¹Early versions of the model are described in Gotz and McCall (1979, 1980). Gotz and McCall (1984) describe the final model delivered to the Air Force.

²The DRM has generally been associated with officer retention, as we are using it in the current project. There is nothing in its basic structure, however, that precludes its use for the enlisted force.

and not to the analyst who simply needs a flexible analysis tool. Once the model's parameters have been estimated, producing policy simulations only requires changing input data such as pay levels or promotion rates.

This Note describes the DRM in nontechnical terms, without recourse to mathematical expressions, and compares its structure and versatility with those of the widely used Annualized Cost of Leaving (ACOL) model.³ A principal goal is to improve the general understanding of the DRM. We show that it actually is a rather simple, common-sense depiction of the determinants of retention in an uncertain world. The seemingly simpler ACOL model we show to be a special case of the DRM that contains an important logical inconsistency as it has generally been applied. The ACOL model has properly been used in the past to examine certain retention issues, but specific scenarios are anticipated in the future for which it would be seriously insensitive. In several examples, including one policy simulation performed by Gotz, we demonstrate the DRM's greater flexibility, which follows from its ability to deal with most policy changes through simple alterations in the model inputs, rather than as fundamental structural changes.

The remainder of this Note is organized as follows. Section II develops the DRM step by step through its assumptions, and shows the relationship between it and the ACOL model. Section III illustrates some of the DRM's capabilities by describing how it would deal with four policy-change examples. Comparisons with the ACOL model shed further light on the important differences between it and the DRM. In Sec. IV we discuss the DRM's data requirements, what is involved in the estimation of its parameters, and how its use for projection differs from that of more familiar regression models. We indicate throughout the sometimes different directions the current project is taking from the original Gotz and McCall formulation.

³The term ACOL appears to originate with Warner (1978).

II. THEORY

The DRM is a two-level model of retention behavior. At one level, it models rational economic behavior under uncertainty. The officer is assumed to choose between staying in the service or leaving, in each year after his initial obligation, based on the mathematical *expectation* of the present value of his future income streams under each choice. He faces uncertainty in his future military pay and in how long he will be permitted to stay, because he does not know whether or when he will be promoted to successive grades. He also is uncertain about the magnitudes of random shocks, such as good or bad assignments, a relative's death, or a new commanding officer's arrival, that might affect his future stay/leave decisions. He is assumed, however, to know the statistical distribution of such shocks, measured as their dollar value to him, as well as the probabilities of his promotion in future years. He also knows the value (to him) of the nonmonetary benefits and disamenities of military employment, net of those in civilian employment; this net value is assumed to be constant over time. Gotz and McCall term this net value the individual's "taste" for military service.

At the second level, the model describes the group dynamics through a set of assumptions about individual discount rates, the forms of the statistical distributions of individual tastes, and the random shocks officers face. The model provides a method to estimate the parameters of those distributions. Given the parameter estimates for the distribution of random shocks, the assumption about discount rates, and estimates of the military and civilian income streams officers can expect to face in the future, the model can predict the probability that an individual with any particular level of "taste" for the military will stay in the military in each future year, under any desired scenario specifying military compensation and personnel policies. Given the parameter estimates for the distribution of "tastes" among the officer population, those stay probabilities can be used to generate retention rate predictions by year of service and grade. If the model is

estimated separately for different "community of interest" officer groups, which would be desirable in any event for predictive accuracy, then the retention predictions can be similarly disaggregated by, for example, source of commission, type of military specialty, and, of course, service.

The model may be described as both forward looking and backward looking. It is forward looking because, in common with other "cost-of-leaving" models, it assumes that the individual's stay/leave decisions are based on the *future* environment he expects to face. The model is backward looking because, unlike other models, it identifies the cumulative effects of *past* stay/leave decisions on the composition of the group--in terms of their tastes for military service--that reaches each grade and year of service. It is the backward-looking feature that allows the model to yield plausible predictions in situations where other models cannot--that a bonus at a given year-of-service point, for example, should *reduce* retention rates (though not necessarily the number retained) in subsequent years of service rather than leaving the rates unchanged.

THE INDIVIDUAL DECISION

The DRM is a discrete-time model, with the period the fiscal year. In each period following the end of his initial active duty service obligation, the officer is assumed to be free either to stay in the military or to leave.¹ Thus, we refer to the first period following the initial obligation as the first *decision point*, the second period as the second decision point, and so on. Perhaps the easiest way to understand the structure of the individual decision in the model is to build the model from the bottom up, beginning with a world of complete certainty and gradually adding complexity.

¹This assumption can be modified to account for multi-year obligations occurring after the initial tour; see the bonus example in Sec. III and the discussion of data in Sec. IV. Such a modification would be desirable if the model were applied to the enlisted force.

Certainty

Consider the officer who has just reached his first decision point, and must decide whether to stay in the service or to leave. The first assumption of the model is:

Assumption 1: The officer acts to maximize the present discounted value of his future income. His discount rate is constant over time.²

If the officer leaves in the first period, he gets a known civilian income stream, and he can calculate its present value. If he stays, however, he will have the same decision to make in each subsequent period. To calculate his return to staying, then, he must decide when he will finally leave. For each period in which he stays, he receives military pay at the beginning of the next period, the level of which depends on his years of service (YOS) and grade in that period. We assume for the moment that he knows when he will be promoted to successive grades, and what the YOS/grade pay structure will be, so he knows what his military pay will be in each period. He also knows what the present value of his civilian income will be for a leave decision in each period, including any military retirement or severance payments to which he would be entitled. Finally, he knows the period in which he will be forced to leave.

One way the officer could approach his first-period problem is to compute the present value of his lifetime earnings for each possible future leave point. One or more of those present values would be greater than all others; that is his return to staying in the first period. The difference between his return to staying and the present value of the civilian income stream he gets if he leaves immediately--his return to leaving--is defined as the *cost of leaving* (or *net return to staying*) in the first period. By Assumption 1, he will leave if the cost of leaving is negative.

²Gotz and McCall assume a constant discount rate, but the model could readily accommodate any assumed pattern of changing discount rates. Constraints on borrowing by young officers (or enlisted persons), for example, could be reflected in high discount rates during early years of service.

A simpler approach than comparing the returns under all possible leave points at once is to examine the various decision points one at a time, asking in each case: If I were still in the military at that point, would I be better off staying or leaving, and what would be the present value of my earnings for the better choice? Answering this question requires knowing the return to the best choice in the next period; thus, the calculation works backward from the last decision point before mandatory retirement. This is the *dynamic programming* approach used by Gotz and McCall. It yields the same cost of leaving as the first approach, and as we add complexity to the problem (below) it offers greater and greater advantages in terms of fewer computations required than the total enumeration described in the previous paragraph.

The first complication we introduce into this simple model is to recognize that the individual's beliefs about the future could change over time. In the example above, the leave point on which the officer decides when he first considers the problem is the exact point at which he later leaves. If his beliefs about military and civilian pay can change, however, this will not be the case. He will have to examine the problem each period, basing his decision on the information he has at that point. We might suppose that the officer's decision would be affected by the possibility that his beliefs about the future environment could change, but Gotz and McCall assume that this is not the case:

Assumption 2: The officer acts in each period as if the future environment will match his current expectations about it.

The officer can learn from his mistakes, adjusting his beliefs about the 1986 military pay raise, for example, after observing the 1985 raise. We have placed this assumption under the Certainty heading, however, to indicate that the officer never expects his beliefs about the future to change until they actually do so.

Uncertain Income

Admitting uncertainty to the model requires some assumption about how officers react to it. For the DRM this is given by:

Assumption 3: *Officers are risk neutral.*

To say someone is risk neutral is to say that he or she will accept an actuarially fair bet. In the current context, it means that officers value uncertain future prospects at their mathematical expectation. For example, if an officer had a 25 percent chance of receiving a civilian income of \$30,000 per year if he left tomorrow, and a 75 percent chance of receiving an income of \$40,000, he would act the same as if his prospective civilian income were $(0.25 \times \$30,000) + (0.75 \times \$40,000) = \$37,500$. This assumption thus implies, in combination with Assumption 1, that officers act to maximize the *expected* present discounted value of their future income streams.³

Income uncertainty can enter in three ways: (1) uncertainty about the pay that will be associated with particular YOS/grade categories in the future; (2) uncertainty about present and future returns to leaving; and (3) uncertainty about whether, and when, the officer will be promoted to successive grades and, if he currently holds a reserve commission, will be offered a regular commission. The first two forms are dealt with by:

Assumption 4: *The officer knows the expected values of his potential civilian earnings for a leave decision in each period, and of future military pay by YOS and grade.*

Uncertainty about future military career paths enters the model more explicitly.

³It is possible to formulate the model to allow for risk aversion, the tendency for individuals to require a positive expected return before undertaking a gamble. For an example, see Gotz and McCall (1979, 1983).

Assumption 5: *As long as he remains in the military, the officer moves among states according to transition probabilities that he knows.*

States are defined by grade, year of service in which promotion to that grade took place (promotion timing), and type of commission (regular or reserve),⁴ with one additional state denoting civilian status. The set of *transition probabilities* gives the probabilities that an officer in any one state at the beginning of a given period will be in each of the other states when the next period begins. These probabilities are determined by the service's policies with respect to promotion, regular commissioning, and high year of tenure. The distinction of states by promotion timing reflects the policy of basing promotion on the number of years in the previous grade, rather than on total years of service, and the likelihood that officers promoted early to one grade will have a better than average chance of being promoted to subsequent grades. Many of the transition probabilities are zero; a one-period transition from grade O-4 to O-6 is virtually impossible, demotions are very rare and assumed not to occur, and certain obvious restrictions apply to movements among promotion timing groups. Most transition probabilities to the civilian state, which denote involuntary separation/retirement, are either zero or one, reflecting high year of tenure policies.⁵

⁴Officers holding reserve commissions may serve on active duty for extended periods, in some cases twenty years or more. In general, however, reserve officers face more stringent high-year-of-tenure policies; Navy Lieutenants (O-3), for example, must leave after being twice passed over for promotion to Lieutenant Commander. Only service academy graduates are assured regular commissions. Officers from other commissioning sources (ROTC, OCS/OTS, etc.) may be offered regular commissions at the service's discretion, or if given a reserve commission may apply for regular force integration (also called augmentation) at any time.

⁵Exceptions to established tenure policies for reserve officers are sometimes granted, leading to involuntary separation probabilities for these officers that may be less than one but greater than zero.

The set of states and associated transition probabilities constitutes the future military environment facing the officer. Thus, Assumption 2 now means that the officer does not expect his beliefs about transition probabilities to change. Note, however, that these are *conditional* probabilities; the Navy Lieutenant recognizes that his chances of being promoted to Commander depend on whether and when he is promoted to Lieutenant Commander.

Career path uncertainty means that the officer will have to evaluate, instead of the one known path, many possible alternative military paths. For each of these he can calculate an expected return to staying, and because he knows the transition probabilities he can attach a probability to each possible path. This allows him to take the expectation across the paths, yielding a single number for the expected return to staying and, thus, for the cost of leaving in the current period. If that cost of leaving is positive, he will stay at least one more period. Future costs of leaving are path dependent, however, so the officer looking ahead cannot predict the exact period in which he will finally leave, only the probability that he will leave in each future period.⁶

Taste for Military Service

All jobs have a variety of nonmonetary features associated with them: working conditions, location, fringe benefits, transfer (rotation) policies, type of housing available, and so on. Gotz and McCall assume that every individual perceives a monetary equivalent for these nonmonetary features of alternative jobs. By monetary equivalent

⁶The structure of the individual decision model thus far is essentially that developed by Gotz and McCall in an unpublished paper ("The Retirement Decision: A Numerical Analysis of a Dynamic Retirement Model," March 1977). It is also described in Gotz and McCall (1979, 1983). They estimated the costs of leaving for officers in various states and years of service and used the results to analyze the effects of alternative retirement system changes. The analysis was rather cumbersome, however, because it involved examining a large number of YOS/grade cells. Chipman and Mumm (1978) and Warner (1978) adopted the simpler approach of regressing retention rates on the calculated costs of leaving for the various YOS/grade cells, taking those retention rates from a single year. Warner called this the *Present Value Cost of Leaving* (PVCOL) model.

we mean the dollar amount that the individual would accept in lieu of the nonmonetary features, and be left equally satisfied with the job.⁷ The annual value an officer places on the nonmonetary features associated with his military service, less the annual value he places on those features he expects in civilian employment, Gotz and McCall term his "taste" for military service (the parameter γ in their mathematical formulation). They assume:

Assumption 6: The annual value of an officer's "taste" for military service is known to him and is constant over time.

Tastes enter the model in a straightforward way. Each period in which he is in the military the officer now receives two amounts, his military pay and the value of his taste parameter. Because the taste value is a *net* amount, it may be either positive or negative.⁸ If it is negative, the individual will only stay in the military if the monetary returns there are greater than in civilian employment. Whether positive or negative, the expected discounted value of the individual's taste is added to his monetary return to staying in each period, yielding a cost of leaving that differs from the purely monetary cost described above. More important, with the introduction of individual-specific tastes two individuals facing the same monetary returns may have different costs of leaving.

⁷An extensive literature deals with the notion that jobs with poor working conditions, fringe benefits, etc., must offer higher wages than jobs with better nonmonetary features if they are to attract equally qualified workers. Brown (1980) summarizes the results of a number of attempts to measure these "compensating wage differentials" for specific job features. Several studies concerned with the value of fringe benefits appear in Triplett (1983); see especially the papers by Smeeding, Smith and Ehrenberg, and Leibowitz. Also in the Triplett volume, Johnson asks whether inter-area wage differences in the United States can be explained by compensating variation in wages to offset area differences in price levels and nonmonetary attributes. In general, these studies are concerned with the *average* value across individuals of nonmonetary job features in various jobs. Here, our interest is more with the range, or *distribution*, of values placed by different individuals facing similar features.

⁸The model could as easily be formulated with the net taste value being something the individual receives when he is *not* in the military--his net *distaste* for military service. This is the approach of the ACOL model.

The addition of tastes allows the model to explain why some officers with negative costs of leaving, as calculated without accounting for tastes, will nonetheless remain in the service, while others whose simple monetary costs of leaving are positive will leave. It also frees the model from the prediction that all officers facing the same monetary incentives will make the same decision.

The complication that individual taste for military service adds to the model is that taste is not something that can be directly observed. We cannot simply take available data on civilian and military pay levels, on promotion and regular force augmentation probabilities, and on high year of tenure policies, and using those data calculate a cost of leaving for all officers in similar circumstances. The individual's cost of leaving depends also on his evaluation of the nonmonetary aspects of civilian and military employment, which is embodied in his taste for military service.

Fortunately we can, by observing the decisions of different officers facing the same monetary incentives, draw inferences about the values of their individual tastes. Those officers who stay at the first decision point may be presumed to have higher taste values than those who do not. Allowing for individual differences requires us to estimate the *distribution* of tastes across a cohort of officers. How that is done, and how it leads to retention predictions, we discuss below. One important implication, however, deserves mention here. The model implies that officers who choose to stay in the military do so in part because they have a greater taste for military service than do those who leave; i.e., as a cohort moves through successive years of service, losing officers along the way, the average value of the taste parameters among those who remain will be higher than among those who left. Comparing two groups of officers--one at 15 years of service, the other at 5--we will find that, even if the two groups' monetary costs of leaving were the same, the older group would have a much higher retention rate, and be much less likely to leave in response to reduced military pay, than the younger group. The thinning of the lower tail of the taste distribution (*censoring*), as those officers with lower values leave the service, explains this phenomenon.

It is worth noting before turning to the next point that the theoretical description of the taste parameter in the DRM may not coincide precisely with the combination of factors that are lumped into the taste parameter in estimation. The parameter will capture all individual differences not otherwise measured by the model, regardless of whether they originate in monetary or nonmonetary factors. For example, suppose we were unable to estimate accurately the civilian earnings opportunities of individual officers, instead taking a simple average of civilian white-collar earnings. To the extent that our estimate is incorrect for any individual officer, the error may be captured by the individual's taste parameter. We say *may* be captured because whether it is perfectly captured depends upon whether the assumption of constant taste over time applies as well to this error in measuring civilian earnings. The variety of factors that may be embedded in the estimated taste parameters makes the term "taste" only a useful first approximation, but to the extent that the factors can be captured effectively they make the model much more robust to errors in specification and in estimated earnings than it otherwise would be.

The ACOL Model.⁹ The structure of the individual decision model at this point is equivalent to that of the Annualized Cost of Leaving (ACOL) model developed by Warner (1979). Consider the first decision point, and suppose for the moment that all officers face the same future military prospects; that is, they are in the same grade and there is no regular/reserve distinction. Looking across officers we see a range of costs of leaving--some positive, some negative--varying directly with their tastes for military service. An officer with a high taste value will have a high cost of leaving, one with a low (perhaps negative) taste value will have a low (possibly negative) cost of leaving. The ACOL model asks the question: What is the taste value for the individual who is just on the margin between staying and leaving, the one whose cost of leaving is zero? The negative of that taste value is called *the ACOL* for the particular decision point.¹⁰ If the ACOL is

⁹Just as the DRM has generally been associated with officer retention, so the ACOL model has been used primarily for the enlisted force. There is nothing inherent in either model that restricts its use to one group or the other.

¹⁰As noted in footnote 8, the taste parameter in the typical ACOL

positive and his taste value negative, it means that the military is offering him more money than he would get if he left; if it is negative, it means that only because of nonmonetary factors is he considering staying despite poorer pay in the military than in the civilian economy.

The ACOL methodology consists of estimating the distribution of tastes across individuals. Given that distribution, then for any calculated ACOL value we can deduce what proportion of the taste distribution lies above the negative of the ACOL, and what proportion lies below. The proportion lying above is the predicted retention rate.

The theoretical structure of the ACOL model creates two major problems that may affect the accuracy and consistency of forecasts. First, the model cannot encompass more than the first decision point without the introduction of a logical inconsistency and an ad hoc adjustment. The ad hoc adjustment allows the model to fit historical data, but does so at the cost of making the model's estimated parameters dependent upon the existing structure of compensation and personnel policies. Second, the model tends to reject the possibility that policy changes affecting military earnings in certain year-of-service ranges will alter retention rates in earlier years.

Suppose that the distribution of tastes were estimated from information on first decision points only (assuming that there were enough observations available with sufficient variation in ACOL values to permit this). The calculated ACOL value for the second decision point would be somewhat higher than at the first, suggesting that the critical taste value separating stayers from leavers is lower. With constant tastes for every individual, this means that everyone who *might* leave at the second point has already done so at the first. Yet we do observe people leaving. To explain this, developers of the model have adopted a statistical methodology that implicitly assumes that individuals' tastes constantly change. This introduces the logical inconsistency. The individual is assumed to act in each period as if his taste value will never change, yet it does. The individual is

formulation is actually the individual's net *distaste* for military service--the negative of the taste parameter as defined here. For consistency, we use the DRM definition of the taste parameter throughout this paper.

continually surprised by his new taste value, never learning from his mistakes.¹¹

The assumption of continually changing tastes breaks the link between policies in one period and retention in the next. As an extreme example of the problem this creates, suppose that a bonus of \$50,000 were offered to every officer who stayed one year longer than his minimum obligation. Virtually all would stay the additional year. What would the retention rate be in the second year? The ACOL model predicts that it would be the same as if there had been no bonus. But those who remained for the additional year only because of the bonus would have little incentive to stay longer.¹² The retention rate at the second decision point should drop, falling so far, perhaps, that only about the same number as before would still be in the service three years after the bonus.

Ignoring the persistence of tastes can create a further problem: a tendency to overestimate the sensitivity of retention rates to military pay. As noted above, persistent tastes imply that the average level of tastes should rise with years of service, as officers with lower tastes tend to leave and those with higher tastes to stay. Retention rates should rise with years of service, therefore, even if financial incentives for staying don't change. Because military pay increases with years of service, however, ignoring the persistence of tastes can lead to attributing the rising retention rates to pay rather than tastes. The greater the dispersion of tastes, the faster will be the rise in retention rates and the more will the retention effects of pay be overstated.¹³

¹¹Warner (1981) develops an ACOL model with persistent tastes, but notes that "for more than two periods the retention equation becomes messy indeed" (p. 12). The two-period model might be useful for predicting first- and second-term reenlistment behavior in the enlisted force, but Warner does not attempt to estimate it. However, it would share the logical inconsistency of the standard ACOL model. In effect, this persistent-taste ACOL model allows the individual's tastes in successive periods to be *correlated* but not necessarily identical, yet he is still assumed to act as if his taste value is constant.

¹²We assume that the payment of the bonus will not by itself lead these officers to expect additional bonuses in future years of service. Neither the ACOL model nor the DRM currently offers a convenient way to estimate the parameters of an expectations-formation process.

¹³Gotz and McCall (1984) estimated ACOL models for Air Force

The ad hoc adjustment is an attempt to return to some persistence of tastes. Users of the ACOL model typically add a year-of-service term when they estimate the model. That is, retention rates are presumed to change with the year of service independently of the influence of changing ACOLs. In effect, this captures the upward shift in the taste distribution for a cohort as lower-taste individuals leave. This pattern of *censoring*, however, depends upon how many individuals leave in each period or, by extension, on the *structure* of military and civilian compensation and of military personnel policies. The YOS term captures the particular pattern in the data used to estimate the model,¹⁴ making the model's parameters dependent on the structure of compensation and personnel policies that gave rise to those data. Anything that changes retention rates, such as the bonus discussed above, will also change the pattern of censoring of the taste distribution. Thus, the more retention rates today differ from those in the estimation period, the worse will the YOS term fit the censoring pattern and the worse will the ACOL model predict the effects of policy changes.¹⁵ In defense of the ACOL model, however, we must note that the practical importance of this problem is uncertain. It may be that in most applications of the model the departures from historical retention

officers using the same data as those used in estimating the DRM. They found that the groups having the greatest dispersion of tastes (and, thereby, the least sensitivity to pay) in the DRM had the greatest sensitivity to pay in the ACOL model.

¹⁴Often, the data for estimating the model are taken from a single cross section. Unless the world is in a steady state, this will mean that the observations for successive YOS points will reflect the differing censoring patterns in successive cohorts. This problem is alleviated somewhat when averages over several years are used, but any persistent trend over those years, such as a steady growth in retention, would still cause problems. A reliance on cross-sectional data, rather than longitudinal, makes the ACOL model really applicable only in a steady-state world.

¹⁵Even if the model starts with the correct pattern of censoring of the taste distribution, it cannot correctly predict the out-year effects of a policy change today. A pay increase, for example, will induce more individuals in YOS-5 to stay, resulting in lower average tastes at YOS-6 next year than was true previously. Thus, the change in retention among YOS-6 individuals should be different in the first year of the pay increase than in the following year. The ACOL model, however, cannot predict this. Warner (1981) recognizes this problem, and describes an approximate adjustment to correct for it.

patterns will be small, or the effects of an altered pattern of censoring minor compared to the direct effects of the policy change that induces them. These are empirical questions that are beyond the scope of this paper.

The second problem arises most strongly as the ACOL model is generally applied, with only a single average military path assumed for every officer. With this simplification, the calculated ACOL for any particular decision point reflects a specific horizon, the planned leave point for the marginal individual.¹⁶ Changes in earnings beyond that horizon generally do not affect the ACOL value, and so cannot change the model's retention predictions for earlier decision points.¹⁷ For officers, the ACOL horizon at the early decision points would probably be 20 years of service; if he stays at all, the marginal officer would plan to stay until he reaches retirement eligibility (but no longer). Thus, the ACOL model would predict no change in early retention if, for example, officers' retirement pay were switched to a flat 50 percent of basic pay for everyone with 20 or more years of service. For enlisted personnel, the ACOL horizon at the first reenlistment point tends to be either four years away (the length of one additional term) or 20 years of service, depending on the assumptions the analyst makes about discount rates and civilian earnings.¹⁸ With the shorter horizon, the

¹⁶The model can be formulated with uncertain future military grades. This yields a set of possible horizons for each decision point, with a probability attached to each reflecting the probability of the military career path giving rise to the horizon. That is, the officer has some probability of being promoted to grade O-6 before reaching 20 years of service. Looking at that possibility from his first decision point, he might determine that his optimal leave point, if he did receive the early promotion, would be 26 years. That would fix one horizon, occurring with a probability equal to his chances of early promotion. Another horizon would be fixed by the due-course promotion path, which would have its own probability. The ACOL for the first decision point would be a blend of the ACOLs for the two paths (not quite a probability-weighted average), and the predicted retention rate could be affected by policy changes affecting returns within either horizon.

¹⁷Only an *increase* in military earnings (or decrease in potential civilian earnings) large enough to move the horizon outward can have any effect.

¹⁸See Warner and Goldberg (1984), p. 31. Only where no reenlistment bonuses are paid, and assuming a discount rate of less than 10 percent, could they find the return calculated over a horizon

model would predict no change in enlisted first-term reenlistment rates if military pay were cut in the senior grades (E-7 and above, say). Indeed, outright elimination of the grades that are rarely reached within the ACOL horizon--O-6 and above for officers, E-7 and above for enlisted personnel--would be predicted to have no effect on early retention. Although again an extreme example, this highlights some of the problems that arise in using the ACOL model. The problems do not make the model indefensible, but do seriously limit the classes of policy changes that the model can address.

Random Shocks

The final element of the individual decision portion of the DRM is the random shocks that are presumed to affect individual officers; the "transient disturbances" in the terminology of Gotz and McCall, represented in their mathematical formulation by the parameter ε . They conceive of these shocks as nonmonetary in nature, giving the example of a family illness leading an officer to leave the service in order to manage the family business. Such nonmonetary shocks are assumed to have monetary equivalents, but as with the taste parameter there is no presumption that two individuals faced with the same shock--a transfer to some remote, unpleasant base, for example--would attach the same monetary value to it.

The officer cannot know the exact value of future shocks, but he is assumed to know something about them.

Assumption 7: The monetary values of random shocks affecting officers' stay/leave decisions are identically distributed random variables, independent over time and with a distribution known to each officer.¹⁹

encompassing YOS-20 to exceed the return over the horizon of a four-year reenlistment. This circumstance may have changed in recent years with the increased use of reenlistment bonuses at the second and even third reenlistment points.

¹⁹The assumption that shocks are uncorrelated across periods may seem unreasonable; a bad assignment, for example, may persist for several years. Ignoring the persistence of some shocks over two or

In each period the officer knows the value of his current shock; if he stays he receives that value, if he leaves he does not.

The random shocks in the DRM are not simply statistical error terms in a regression, as they are in the ACOL model. Rather, they form an integral part of the environment facing the individual; they can be viewed as random fluctuations in the factors determining his taste for military service. An officer who finds the military rank structure too rigid for his liking might stay in well past the end of his initial obligation if he received assignments involving little contact with other military personnel, while another who loves all things military might yet leave if he were assigned to a radio listening post in Greenland. This is why we have talked above about a *censoring* of the taste distribution, rather than a *truncation*; at any YOS point, some officers will remain who have very low taste for the military, and others who have high taste will already have left. Further differentiating this model from ones in which the shocks are lumped into an error term, it explicitly recognizes that an officer is more likely to stay today if he knows that he can avoid large negative shocks by leaving the military in the future.

The DRM description of the environment facing the officer and how he reacts to it is stylized, but it captures the chief factors relevant to decisionmaking. It recognizes that individuals are different--that some love the military, some hate it, and some consider it just another job. It explicitly treats the officer's natural uncertainty about the future--about how quickly he will advance in his military career and whether some event will cause him to change his stay/leave plans. Combined with the distributional assumptions discussed below, the description of the individual's decision provides a powerful model for predicting officer retention rates under a wide variety of policy environments. We discuss some of the model's possible applications in Sec. III.

three periods, however, is not likely to affect significantly the model's predictive ability.

AGGREGATE ASSUMPTIONS

The model so far describes how rational individual officers with assumed preferences and beliefs would act in the face of the uncertainty inherent in the environment they face. To turn this into an estimable model of aggregate retention rates Gotz and McCall make further assumptions about: (1) the commonality of preferences and beliefs across individuals, (2) the form of the distribution of random shocks, and (3) the form of the distribution of "tastes" across individuals.

Assumption 8: All officers have the same rate of time preference, and thus apply the same discount rate in calculating present values.

It is possible to formulate the model with heterogeneous discount rates, but it is not clear that this added complexity would improve the model significantly. Some aspects of such heterogeneity are presumably captured in the current formulation's "taste" parameter.

Assumption 9: Officers have identical beliefs about the distribution of the random shocks each faces.

When combined with Assumption 7, Assumption 9 implies that the true distribution is the same for all officers.

The last set of assumptions is fairly technical in nature, specifying the functional forms of two distributions, and may be skipped without loss of continuity.

Assumption 10: The random shocks are distributed normally, with mean zero and variance σ^2 , and are independent across individuals.

This assumption pins down a probability distribution that has hitherto been described in general terms, restricting it to a family of distributions--the normal--that is characterized by two parameters, a mean and a variance. The variance will be estimated, the mean is assumed to be zero because the model does not permit it to be identified

separately from the mean of the taste distribution.²⁰ Normality is assumed because it is plausible. The normal distribution is symmetric (positive and negative shocks of equal absolute value are equally likely), and although it does not rule out any size of shock, it is "thin tailed" (extremely large shocks are rare). It entails only one parameter to be estimated, an important advantage over some more-flexible functional forms.

Assumption 11: *The values of the "taste" parameters of officers at their initial decision points follow an extreme value distribution for maxima, with location parameter θ and scale parameter ω . The values of θ and ω are the same for all cohorts.*²¹

The extreme value distribution is skewed to the right; that is, values much greater than the mean are more likely than are values much less than the mean (see Fig. 1). This seems plausible, for individuals with a great distaste for the military are not likely to have entered in the first place.

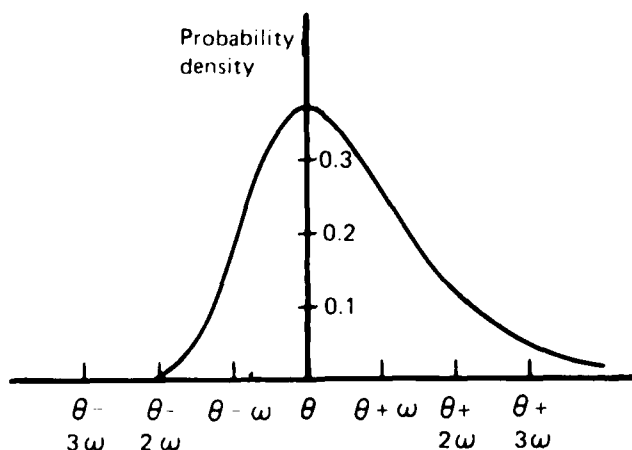


Fig. 1 — Probability density of the extreme value function for maxima

²⁰The average value of the random shocks may be thought of as one of the nonmonetary features of military employment that determine tastes. Lumping the average value of the random shocks into the average taste value creates no loss of explanatory power.

²¹This assumption of identical taste distributions across cohorts can be relaxed. In the course of our model estimation, we plan to determine whether any change in the assumption is warranted.

One final assumption solves a problem that Gotz and McCall encountered in some of their early empirical work with the model.

Assumption 12: *The probability that an officer will hold a regular commission (rather than a reserve commission) at his first decision point is dependent upon the value of his "taste" parameter.*

The formal specification of this relationship between taste and type of commission is presented in Gotz and McCall (1984), and will not be repeated here. It introduces a new "selectivity" parameter, α , which expresses the strength of the relationship. In principle, the relationship could be either positive or negative; that is, an officer with a high taste for the military could be either more or less likely to be offered a regular commission than an officer with low taste. Gotz and McCall argue, however, and find empirical support for their belief, that officers with high taste values will be more likely to hold regular commissions, presumably because those who most want to stay in the military will tend to perform better than those who want to leave. Their estimates of α imply that the average taste among officers who receive regular commissions before the first decision point exceeds that among officers who do not.²²

To summarize, the aggregate assumptions lead to the need to estimate at most five parameters.

1. The *discount rate* describes the relative importance to individuals of income in the present versus income in the future.

²²The variance of tastes among officers who have stayed past their first decision points will be much lower than among officers who have not yet reached that point. Thus, although the service may continue to offer regular commissions predominantly to higher-taste officers, the effect of accounting for this will be less. For computational convenience, Gotz and McCall assume that the relationship disappears after the first decision point.

2. The *variance of the random shock distribution* (σ^2) measures the "spread" of these transient disturbances, in effect describing how large is a "large" shock.
3. The *location parameter of the taste distribution* (θ) measures the mode of this distribution (the point of greatest density). Because the assumed form of the distribution is skewed to the right, the mode is smaller than the mean.
4. The *scale parameter of the taste distribution* (ω) measures the spread of this distribution (the standard deviation is equal to ω times 1.28). It also determines how much greater than θ is the mean (equal to $\theta + 0.577 \omega$).
5. The *selectivity parameter* (α) measures the strength of the relationship between an individual's taste value and the probability that he will hold a regular commission when he reaches his first decision point. Note that for officers graduated from one of the service academies, this parameter is irrelevant; all are given regular commissions.

The estimation of these parameters, and of the earnings and transition probabilities required by the model, is discussed in Sec. IV.

III. APPLICATIONS

The DRM is capable of predicting the effects of a broad range of compensation and personnel policy options, including many for which there are no historical analogues. It can do so because its parameter estimates describe, not the average response of a group of officers to specific external stimuli, but rather the *preferences* of the officers themselves and a key aspect of the environment they face. It models an officer's *decision process*, transforming policies affecting his future in the military into current monetary equivalents in the way that a rational individual satisfying certain assumptions would. As a result, the parameter estimates are less dependent on the specific structure of policies in effect during the period from which they are obtained than is true for previous models. As a further result, changes in those policies can be readily accommodated, so long as their effects on the officer can be described in terms of monetary returns or the probabilities of different returns.

The versatility of the model derives from three features. First, it incorporates a very rich picture of officers' military futures, reflecting the divergent paths they follow when some are promoted and others are not, when some are offered regular commissions and others are not, and as a result of the first two, when some must leave and others are allowed to continue. The major personnel policies determine the probabilities with which any officer will follow the various possible paths, and compensation policies determine the returns along each. Second, it recognizes that individuals differ in persistent ways, which it describes in terms of the values they place on nonmonetary aspects of life in the military and in civilian employment. The persistence of these differences gives rise to the model's "backward-looking" retention rates; the retention rate for the group of officers at any particular decision point depends upon who is there to be making a decision--what the distribution of tastes is among them--which in turn depends upon the history of compensation and personnel policies preceding the point. Third, the model reflects the possibility that unforeseeable

circumstances may cause any officer to leave no matter how strong his attachment to the military, or to stay beyond his initial commitment no matter how strong his distaste for military life. Thus, it allows retention rates to be influenced not only by policy changes affecting the marginal officer--the one on the borderline between staying and leaving--but also by changes that affect only some infra-marginal group.

We illustrate the importance of the three features in giving the model its versatility as a policy simulation tool with four examples: (1) a change in the involuntary separation policy for officers who fail to be promoted to O-4; (2) a temporary, two-year postponement of active duty pay increases; (3) a series of bonuses for officers in the fifth through twelfth years of service, alternatively with and without an active duty service obligation; and (4) a change in the retirement system similar to one recently proposed by the Fifth Quadrennial Review of Military Compensation (QRMC). To provide a basis for comparison, we also indicate the nature of the predictions that the ACOL model would make in each case. As will be seen, the ACOL model is unable to deal fully with any of the four policy changes, although it might produce usable partial predictions if it were carefully used.

UP-OR-OUT POLICY

The four services currently have some latitude in their treatment of officers who are twice passed over for promotion to O-4, so for simplicity we will assume for our baseline that the policy is to require such officers to leave within one year after the second failure.¹ We assume for concreteness that this separation would occur in the twelfth year of service. The change we examine is allowing the passed-over officers to stay until 20 years of service. Such a change might be instituted as part of a general movement outward in high years of tenure, but for this example we consider its effects in isolation.

One of the future paths the officer examines in the DRM involves failure of promotion to O-4. In the baseline case, this path is terminated in the twelfth year of service by a transition to the civilian state occurring with probability one (involuntary separation).

¹This is the current policy for Navy officers holding reserve commissions.

To reflect the change, the transition probability would be reduced to zero (or slightly greater if there are other reasons for involuntary separation). The passed-over officer then gains both the opportunity to stay longer and, if he stays to twenty years, retirement benefits (we assume that an earlier departure would still entitle him to any severance payment available in the baseline case). Officers in earlier years of service who have not yet been promoted to O-4 will thus see their returns to staying in the military increased, since the value of being in at the beginning of the twelfth year, along the passed-over path, is increased, and they have some positive probability of being on that path. The model will predict that retention in the earlier years will increase, and will provide predictions of retention rates for passed-over O-3s in the twelfth through twentieth years of service.

The ACOL model could, in principle, accommodate this policy change. As the model typically is implemented, however, both its input data and its output predictions do not differentiate grades.² The observation at each YOS point reflects the *average* grade (as measured by military earnings) of the officers at that point during the estimation period. Because the change in up-or-out timing affects only some of the officers in the 12-plus YOS range (those in grade O-3), it thus appears as a structural change to the model. It would, of course, alter the grade distribution beyond the twelfth year of service. A user of the ACOL model would have to be careful not to allow this change in the grade distribution to be reflected in a reduction in average military earnings in those years. If he did, the model would give the erroneous prediction that retention in the earlier years would *fall*. The model could not, of course, provide separate retention predictions for the passed-over officers; it would implicitly predict that their retention rates would be the same as those of officers who had been promoted to O-4. In practice, however, this would likely be a minor error.

²Gotz and McCall (1984) derive an ACOL model with the same probabilistic depiction of the individual's future movement through grades that is used in the DRM. Warner and Goldberg (1984) apply this structure. Adding the probabilistic depiction, however, deprives the model of much of its apparent simplicity, which seems to be its principal attraction.

POSTPONEMENT OF ACTIVE DUTY PAY INCREASES

On several occasions in recent years, military pay raises were deliberately held below the levels necessary to match increases in civilian-sector pay. The FY82 raise, however, was intended to offset earlier shortfalls. We examine here a similar policy: a two-year pay freeze followed in the third year by a "catch-up" that restores military pay to the same relationship it held to civilian pay before the freeze. That is, military pay increases are *postponed* for two years, and we assume that officers know it is only a postponement, not a permanent reduction in relative pay. We compare this with a baseline policy of matching civilian pay increases every year, and assume for concreteness that the freeze leaves military pay 10 percent behind civilian before the catch-up. The pay freeze will affect officers in every year of service, but for illustration we examine its effects only on the group that is first eligible to leave in the initial year of the freeze, and assume that this is the fifth year of service for these officers.

The DRM and the ACOL model will offer similar predictions for the short-run effect of the freeze. The military earnings reduction in the fifth and sixth years of service, relative to the baseline case, will reduce the returns to staying in the military, and thus reduce the fifth-year retention rate. The reduction in retention will not be as large as if the 10 percent fall in relative pay were permanent, of course, but it should still be sizable because earnings in the near future are more important to individuals than are earnings farther out. The ACOL model will also predict a fall in the retention rate for these officers when they reach their sixth year, but for the DRM the direction of the change is not so clear. The reason for this can be seen when we turn to the seventh year, in which the catch-up restores military pay to its level under the baseline case.

To the ACOL model the two cases look identical in the seventh year, so it will predict identical seventh-year retention rates. The DRM, however, will recognize that the smaller number of officers still present after the freeze will, on average, have a higher level of taste for military service than will those under the baseline case. The group that the freeze induces to leave the military will be composed

disproportionately of officers with low attachment to the military. Thus, the DRM will predict that the retention rate in the seventh year of service will be *higher* for this group that lived through the pay freeze than it would have been had there been no freeze.³

In the sixth year of service, the DRM will account for two conflicting effects: on the one hand, military pay will be lower than under the baseline case, leading to lower retention; on the other hand, the officers who are still in the service after the first year of the freeze will have a stronger attachment to the military than had there been no freeze, and so will be more likely to stay. We would expect the first effect to be the larger, but the net reduction in the sixth-year retention rate would certainly be smaller than the ACOL model would predict. Because the ACOL model does not account for the censoring of the taste distribution as some officers in a given year-group leave, it yields a predicted retention rate for the seventh year that is clearly too low, and a similar though less obvious error for the sixth year. In the DRM, the backward-looking retention rates reflect the history of compensation and personnel policies in earlier years. Although the practical importance of this in the current example cannot be known without empirical investigation, it appears that ACOL model users examining such a policy change would be well advised to use the approximate correction for prior pay changes described by Warner (1981).

OFFICER BONUSES WITH AND WITHOUT OBLIGATIONS

For this example we draw upon simulations performed by Gotz and McCall in the course of their development of the model.⁴ We examine the effects of a series of bonuses given to officers in their fifth through twelfth years of service, alternatively with and without active duty

³Although the retention rate will be higher, the *number retained* past the seventh year will be lower. Those who remain will act about the same as they would have under the baseline case, but some of those who have been lost because of the freeze would also have stayed in the seventh year, leading to the prediction--when the model is integrated with an inventory projection model--that the number of officers present in the eighth year will be smaller.

⁴These results refer to non-rated Air Force officers who received their commissions through ROTC, and were first eligible to leave when they began their fifth years of service in FY83.

service obligations. The amounts and timing of the bonuses, modeled after ones being considered for science and engineering officers when the simulations were performed, appear in Table 1. In the without-obligation case, all officers who decide to stay in the fifth year must remain through the eighth year, receiving the bonuses along the way. If they stay in the ninth year they must remain through the twelfth.⁵

Table 2 shows the simulation results, following a single year group through successive years of service. Thus, if the YOS-5 column refers to FY83, the YOS-6 column refers to FY84, and so on. Looking first at the without-obligation results, we see that the bonuses would elicit a small retention rate increase over the base case in the fifth year of service, when the \$6,000 bonus is paid, and again in the eighth and ninth years. The YOS-8 result is attributable to the anticipation of another \$6,000 bonus in the ninth year. None of the effects is very large because the bonuses themselves are not large; the \$12,000 received over years five through eight is only about six times an officer's monthly basic pay, or about a 12 percent addition to his pay over the four years. The ACOL model would give similar predictions for the without-obligation bonuses.

Table 1

OFFICER BONUS
AMOUNTS AND TIMING

Year of Service	Amount
5, 9	\$6,000
6, 10	2,000
7, 11	2,000
8, 12	2,000

⁵This policy is more akin to the reenlistment-with-bonus practice in the enlisted ranks than it is to any proposal actually made for officer bonuses. We use it to illustrate the negative retention effect of a fixed obligation.

Table 2
EFFECTS OF OFFICER BONUSES
ON YEAR-GROUP VOLUNTARY RETENTION RATES

	Year of Service				
	5	6	7	8	9
Baseline	57	94	99	98	99
Without obligation	59	94	99	100	100
With obligation	56	100	100	100	100

When an active duty service obligation is added, the DRM shows a rather different effect of the bonuses: a fall in retention in the fifth year of service. Why is this? We can see from the without-obligation results that the bonuses alone should elicit a 2 percentage point increase in retention, but the four-year obligation takes something away from these officers: their flexibility to leave in years six through eight. Because they know that something unexpected (a random shock) might cause them to want to leave during those years, they respond to that flexibility loss by staying in smaller numbers. The retention rate at YOS-5 drops one point from its 57 percent level in the baseline case. Knowing that the bonuses should increase the YOS-5 rate by two points, we can conclude that the loss of flexibility reduces the rate by 3 points. Were we to drop the bonuses (i.e., require a "reenlistment" at the fifth year of service) the retention rate would drop from 57 percent to 54 percent.

The ACOL model would predict no change in retention at YOS-5 when the active duty service obligation is added. The annualized cost of leaving used by the model at that point is that of the *marginal* officer. Given almost any feasible pattern of compensation over years 5 through 19, the marginal officer will be planning to stay until he reaches

retirement eligibility at 20 years of service, if he stays at all.⁶ Thus, changes that do not alter his military earnings prior to retirement cannot affect his decision at the fifth year of service. Flexibility--the ability to leave whenever the officer wishes--is not something that is valued. The ACOL model will predict that requiring officers to commit themselves to four years of service if they wish to stay at YOS-5 will not affect the retention rate at that point.

We cannot be sure, of course, that the DRM predictions will be more accurate than those of the ACOL model. The methodology of the DRM does not preclude, however, a finding that officers do not value flexibility--the implicit assumption of the ACOL model. If the variance of the random shocks (σ^2) were estimated to be zero--if, that is, officers acted as if they knew the future with certainty--the simulation results in the with- and without-obligation cases would be identical. In this context, the ACOL model can be viewed as a limiting special case of the DRM. Gotz and McCall (1984) produce estimates of the random shocks facing Air Force officers that are quite far from zero, indicating that ACOL model users should be very cautious in applying it to cases involving changing obligations.

QRMC RETIREMENT PROPOSAL

The Fifth QRMC proposed a number of alternative reforms to the military retirement system. The one its chairman reportedly favored has the following features: (1) the monthly annuity payment is reduced from the current 50 percent of basic pay after 20 years of service to 35 percent, rising with additional years of service to the current 75 percent at 30 years; (2) a lump-sum payment equal to 200 percent of

⁶This point illustrates the theoretical inconsistency in the ACOL model. Although the marginal officer at the fifth year of service plans to stay until 20 years, the model does predict retention rates of less than 100 percent in the sixth and subsequent years, when the monetary costs of leaving are higher. It does this because in each of those years every officer is implicitly assumed to take a new draw from the taste distribution. Were it not for the YOS term included in ACOL regressions, the model would predict a slow rise in retention rates from one year of service to the next, rather than the very rapid initial rise actually observed for officers (see Table 2). The YOS term moves the taste distribution to the right (higher average taste) to account for censoring, but it is clear that this is an ad hoc adjustment, albeit one that may fit historical retention patterns fairly well.

annual basic pay is made at retirement; and (3) individuals with 12 years or more of service remain with the old system, while those with less than 12 years may choose either system.⁷ The second provision seems particularly unlikely to emerge from Congress unscathed--typical payments would exceed \$70,000 at current pay levels--so for our example we assume that the lump-sum payments would be only 100 percent of annual basic pay.⁸

Figure 2 shows the nature of the proposed change in terms of the basic pay percentages of the retirement annuity.⁹ The lump-sum payment has been converted to its annuity equivalent under the assumption of a 7.5 percent discount rate and an infinite lifetime (using a finite lifetime of typical length would not markedly affect the computation). The figure makes it evident that the proposal (as we have modified it) would make retirement in the early years less attractive, and in the later years more attractive.

What does the DRM have to say about the effects of adopting the QRMC proposal? First, it should be clear that, if the taste distribution among officers reaching retirement eligibility is unchanged, the new system would tend to make them stay longer. An unchanged distribution would arise over the next few years if the proposal were adopted immediately (without prior discussion) and included no grandfathering provision (i.e., all officers are covered by the new system). That the long-term effect of the proposed system would be to change the distribution can be seen, however, by considering who, under the actual grandfathering proposal, would choose the current system and who the new.

⁷See "Top Leaders Oppose Retirement Overhaul," *Air Force Times*, 23 January 1984. The proposal also would limit annual cost-of-living adjustments in retirement payments, but for simplicity we ignore this provision.

⁸Without presuming to read the minds of the QRMC staff, we note that at roughly a 7.5 percent real discount rate the 200 percent lump-sum payment would offset the reduced present value of the retirement annuity for officers who retire at 20 years of service.

⁹The line for the QRMC proposal should actually be slightly curved; the rise in the percentage of base pay with years of service is slower in the early years than later. For simplicity of exposition we assume a linear rise.

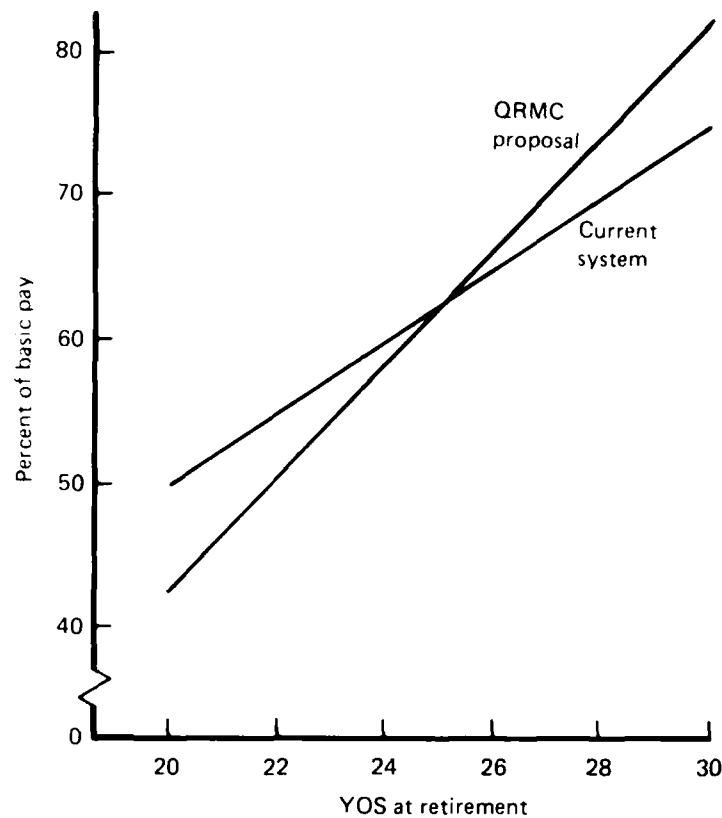


Fig. 2 — Comparison of annuity equivalents for QPMC proposal and the current retirement system

Two factors would affect the choices of officers between the two systems: their tastes for military service and the up-or-out system. We abstract from the second for the moment by assuming that all officers are allowed to remain in until they complete 30 years of service. Those choosing the current system, then, would be those who expect to retire in the early years (20 to 24), whereas officers expecting to stay longer would choose the new system. The first group would consist of officers with relatively low taste for military service; the second of those with relatively high taste. Thus, we might expect that in the long term, when all officers are covered by the new system, the average taste of officers reaching retirement eligibility would be higher, implying that fewer would stay to 20 years or more. Whether this would occur would depend as well on whether the greater benefits available at 25-plus

years of service would induce an offsetting rise in retention, but at reasonable rates of discount such an effect appears unlikely.

Current up-or-out policies do not permit all officers to stay to 30 years of service; in general, passed-over O-3s must leave at 20 years, O-4s at 24, and O-5s at 28. Thus, officers who believe their chances of promotion to O-5 and O-6 are poor will tend to opt for the current retirement system. This will be captured in the model through the timing of promotion to O-4. Officers who receive early promotion ("below the zone" or "deep zone," to use the services' terminologies), will rightly perceive their chances of subsequent promotion to be greater than average, and those who have been passed over will perceive their chances as poor.¹⁰ The richness of the future paths considered in the model allow these differences in subsequent transition probabilities to be accounted for explicitly. The model will predict that "fast-track" officers will tend to choose the new system and "slow-track" the old and, more generally, that the group choosing the new system will be composed disproportionately of officers who perceive themselves as doing well in the military system. In the long run, the QRMC proposal and the current up-or-out system would apparently bring to retirement eligibility fewer of those officers whom the services have judged to be less capable and, by extension, fewer officers in total.

As is probably clear at this point, the DRM is capable of predicting the proportion of officers (with less than 12 years of service) who will choose each system. It will also predict a *one-time* increase in retention; those choosing the current system will not change their retention behavior, but those choosing the new system will do so because it makes them better off, and thus will be more likely to stay than otherwise. For yeargroups entering the military after the change the model would also yield retention predictions, but although we have suggested above that retention rates would generally fall in this specific illustration, a full simulation would be required to confirm this conclusion.

¹⁰In general, both early promotions and promotions to O-6 tend to be concentrated in certain specialties--pilots in the Air Force, for example. Officers in these favored specialties would seem to be more likely than others to choose the new system.

The ACOL model would be of more limited usefulness for predicting the consequences of this proposed retirement system change. It could probably make fairly accurate predictions of the reduction in pre-retirement retention rates, although the improved retention of those infra-marginal officers who expect to stay well beyond 20 years, if they stay at all, might make the ACOL model's predictions somewhat too pessimistic. The model would not predict the one-time increase in retention rates, but this effect might also be small. Where the model would prove least useful is in examining the more subtle effects of the change: how many of those officers allowed to choose would elect the new system, how the character of the officer groups reaching retirement eligibility would change, and the extent of the shift toward later retirement. It would be possible to examine the last effect with the aid of an ACOL model, but there is a good possibility that the model would underpredict the long-term shift because of its inability to account for the altered taste distribution at the 20 YOS point.¹¹ The model is better suited to dealing with retirement-system changes that are limited to simpler upward or downward shifts in the annuity-by-YOS curve.

SUMMARY

Our examples were obviously chosen to highlight some of the important differences between the DRM and the ACOL model. In doing so, however, they have also demonstrated that the greater complexity of the DRM provides some very real benefits. The model provides plausible predictions in situations where others cannot because it relies on a description of the individual's decision process, and of the environment he faces, that is both rich and intuitively appealing. It incorporates many of the important policy parameters, so that changes in those parameters need not be treated as structural changes to the model, and

¹¹We should note that the proposal actually offered by the QRMC, which included lump-sum payments equal to 200 percent of basic pay, would probably have a less marked effect on the taste distribution than would the alternative we examine here. Thus, the QRMC analysis of the proposal, although based on an ACOL model, cannot be strongly criticized in the area of shifts toward later retirement.

it deals with all officers, not only with those on the margin at each decision point. It does not describe all of reality, of course--rotation policies, for example, do not appear--but it seems to capture the most important elements. In short, the DRM offers sufficient advantages to justify at least some additional effort, over that required for an ACOL, in its estimation and implementation.

IV. DATA, ESTIMATION, AND PREDICTION

The versatility of the DRM does not come without cost; the model requires more data than other models, and a more complex estimation procedure. A potential user of the model should be aware of these requirements, although they will not make the DRM difficult to use once its parameters have been estimated. This section describes the data required for both estimation and prediction, discussing in the process the data being assembled for the current study. It then gives a non-technical explanation of the maximum-likelihood technique used in estimating the model's parameters. The section ends with a discussion of the differences between prediction with the DRM and with simpler regression models such as the ACOL.

DATA

The data required by the model fall into two general categories: data on officers' retention decisions at their first and subsequent decision points and data on officers' earnings opportunities both in the military and in the civilian sector. The data on earnings opportunities must cover not only the period of the retention decisions being considered, but also well into the future. They are further divided into: (1) transition probabilities among military states; (2) military earnings by grade, year of service, and community (for special pays); and (3) civilian earnings streams for leave decisions at various YOS points.

Retention Decisions

The basic observational element Gotz and McCall call the *event*. An event describes the relevant elements of an officer's military history, including the following:

1. The year in which he was first eligible to leave.

2. His community (military specialty or group of specialties).
3. His source of commission (Academy, ROTC, OCS/OTS).
4. The year of service in which he was first eligible to leave.
5. His stay/leave decision in the last year observed.
6. The sequence of military states he occupied.¹

An event could be constructed, for example, from the following information about officer X: He is a pilot, an ROTC graduate, held a reserve commission in FY78 when he was in his fifth year of service and first eligible to leave, was a Captain (O-3) in that year, was given a regular commission in the following year, was promoted to Major in due course two years later (ninth year of service), and was still in the service at the end of the next year. This officer's history ends when the data run out at the end of FY83; others will end because of a leave decision.

It is not necessary (or possible) to observe every officer's complete history. Most of the information that the events provide comes from the first few years following the end of the initial obligation, when most decisions to leave the service take place.² Thus, we would not necessarily want to exploit a long period of data by constructing equally long events for the officers who make their first stay decisions in the first year we observe. Rather, events can probably be limited to covering five years. Some, toward the end of the data period, may cover only one year, but they cannot all be this short or the model's parameters cannot be identified.

Although long events are not required, a long data period offers two advantages. First, it provides variation in the environment faced by officers, which should yield variation in retention rates. Having several years of data, with some years having high retention at the first decision point and others low, provides more information about the

¹Recall that states are defined by grade, component (regular or reserve), and timing of promotion to the current grade.

²Many leave decisions also take place in the retirement-eligible years. We are exploring ways to take advantage of the information those decisions provide.

shape of the taste distribution (the scale parameter ω) than if only one or two years were available or if retention rates never changed. Second, a long period allows a test of one of the maintained hypotheses of the model, that the distribution of tastes among officers at their first decision points (and with the same community and source of commission) is the same in every year. If that hypothesis appears to be wrong, we could explore possible determinants of the distribution, such as the sizes or demographic characteristics of the different cohorts.

The requirement for longitudinal data (successive observations on the same individual) comes basically from the necessity for identifying two distributions: the distribution of tastes and the distribution of random shocks. One may think of the first decision point as providing information primarily about the location and spread of the taste distribution. A year in which monetary costs of leaving are high (and thus retention is high) identifies a point low on the taste distribution, and one in which the costs are low picks out a point high on the distribution. The second (and subsequent) decision point provides information about the variance of random shocks. Supposing for the moment that the monetary costs of leaving were the same at the second decision point as at the first, then those who choose to leave at the second point would either be avoiding a negative shock in that period or have experienced a positive shock in the previous period. The monetary costs of leaving will not stay constant, of course, but this variation provides additional information about the sizes of the shocks that must have occurred to lead to the observed leave decisions at the second and subsequent points.³

A key requirement of the officer history data is that they indicate the year of service in which the officer was first eligible to leave. If we assume that all the officers in a particular cohort could leave after four years of service, when in fact half had commitments of five years, we would infer erroneously from the large number of leavers in

³It might appear that information on only the first decision points of successive yeargroups would permit the identification of the taste distribution parameters but not that of random shocks. This is not the case. One could never be sure whether changes in retention rates as the costs of leaving changed were due to a dispersion of tastes or to a dispersion of random shocks.

the sixth year that random shocks must be very large and the variance of tastes very small. A difficult problem is presented by individuals whose initial obligation is extended for some reason: Should their stay decisions in the extension period be viewed as voluntary, or should the end of the period be treated as the first decision point? Incurring an additional obligation--by accepting additional training, for example--clearly is a voluntary choice, but it may be so much a part of a normal career progression that any officer with even the slightest expectation of staying to ten years or more would have to incur the obligation. Gotz and McCall took the approach of treating the first change in an officer's end of obligation date as being effectively involuntary, but any subsequent change as voluntary. That is, only the first change alters the initial decision point for the officer in question.

The data for constructing events are being taken in the current study from master and loss files maintained by the Defense Manpower Data Center (DMDC), which reliably cover the period FY73 to the present. We chose this source, rather than requesting separate submissions from each service, to minimize "learning-curve" time and to enable much of the data processing to be performed at DMDC. The data processing programs developed in this study will be transferred to DMDC for their future use, thus simplifying subsequent updating of the data and reestimation of the model's parameters when that proves necessary. One problem that has arisen--a lack of specialty information on Navy officers in the DMDC files--has been resolved through the submission by the Naval Personnel Research and Development Center (NPRDC) of a file containing Navy community designators for all officers on active duty during the relevant period.⁴

The ACOL model. In contrast to the DRM, the ACOL model typically is estimated from a single cross section. The basic observational element is the retention rate at a particular year-of-service point. Because the model does not admit of any connection

⁴A second problem came to light as this was being written: the Navy does not record extensions of the initial active duty service obligation in any of its automated files. We have developed a method, however, for treating stay decisions in the first two years following the end of the initial obligation as if they reflected, with some probabilities (to be estimated), additional commitments incurred essentially involuntarily.

between the retention rate at YOS-5 (say) in one year and the rate at YOS-6 in the following year, it does not require longitudinal data. Additional years of data, if they are used, simply provide variation in the costs of leaving. The simplicity of the historical data required by the ACOL model appears to be the model's principal attraction to policy analysts.

Transition Probabilities

The information required to construct the transition probability matrices (one matrix for each community/service) can be obtained from the same DMDC files as the events, supplemented by information supplied by the services. These historical data are not perfect, however; we would prefer to know the probabilities that officers *believed* they would face in the future at the times they made their decisions. Thus, we implicitly assume that officers form their expectations about the future in the same way that we derive the probabilities: by observing promotion and regular force integration rates in the recent past. Such information, particularly the promotion rates, would be available to the officer through articles regularly published in the service *Times*.

The promotion probabilities that the model requires are not simply the announced promotion *opportunities*, the percentages of a year group expected to be promoted. The timing of promotions is important as well, both because it determines when the officer's military earnings change and because his chances of promotion to one grade are allowed to depend upon when he was promoted to the previous grade. This dependence derives from two sources: (1) an officer promoted to one grade a year early ("below the zone") will be in the primary zone for promotion to the subsequent grade one year earlier than his contemporaries who received due-course promotions; and (2) early promotion to one grade is a signal, both to the particular officer and to an outside observer, that the officer is viewed by the service as more capable than his due-course contemporaries, and thus is more likely than they to be promoted to higher grades. Of course, whether the second source of dependence is actually important must be determined from the data (Gotz and McCall did find it for the Air Force), but it is explicitly allowed for in the definition of states by grade, component, and the year of service in which the officer received his last promotion.

Promotion and regular force integration rates vary over time. Thus, a separate transition probability matrix is required for each year of the estimation and projection periods. As a starting point, the matrices for the projection years may be assumed to be the same as in the last year of the estimation period. Eventually, however, it would be desirable to derive the future matrices from announced service plans, or to consider possible alternative sets of policies. An inventory projection model would be needed to do this; transforming a policy lever such as a promotion opportunity into a set of promotion probabilities incorporating timing requires information on the numbers of officers in the relevant years of service. A service that manages by yeargroup and maintains a constant opportunity, for example, must alter promotion timing as larger and smaller yeargroups move through. This process can be mirrored in the inventory projection model, which would then supply the transition probabilities to the retention model.

The ACOL model. In place of the transition probabilities of the DRM, the ACOL model as it is usually applied substitutes a single "average" path for all officers.⁵ This average is really a set of average earnings over the officers in various grades at each year-of-service point. Determining the average earnings at any point requires combining military pay information (below) with the proportions of officers in each grade, which typically are derived from a single cross section. This procedure's simplicity is another attraction of the ACOL model, but it limits the model to examining only policy changes that do not affect promotion rates or timing, high years of tenure, or, more generally, the grade/YOS distribution.

⁵As noted in Sec. II, Warner and Goldberg (1984) calculate ACOL values using a transition probability matrix. Their promotion probabilities are distinguished by length of service rather than time in grade, however, and they do not indicate whether they account for high year of tenure policies.

Military Earnings

Data on military earnings can be obtained from published pay tables. For communities in which officers are eligible to receive special pays (flight pay, submarine pay, etc.), information on average payments by grade and YOS will be obtained from the services. For the future, a reasonable baseline assumption is that military pay will continue to grow in real terms at its average rate over some recent historical period.

Civilian Earnings

Estimating the civilian earnings opportunities of servicemembers is a problem faced by users of any model based on costs of leaving, in which category both the ACOL model and the DRM fall. Gotz and McCall (1984) describe the procedures used in the earlier study. Basically, they used average earnings data for civilians with education levels similar to those of the military officers in question, and assumed a certain pattern of wage penalties reflecting the presumably imperfect transferability of military skills to civilian jobs. For the current effort we are exploring possible improvements in the Gotz and McCall methods, principally involving the use of data on the actual civilian earnings of former officers. These data will come from two sources: (1) the 1977 DoD Retiree Survey and (2) a file assembled for the Fifth Quadrennial Review of Military Compensation that merges the IRS and Social Security earnings of a sample of military separatees (including retirees) with information on their status (grade, specialty, etc.) when they left the military.

There are two drawbacks to using average civilian earnings of both veterans and nonveterans to estimate the potential civilian earnings of those in the military. First, military officers form a nonrandom sample of similarly aged and educated individuals in the general population. Our ability to control for the determinants of earnings is limited to observable characteristics; those who enter the military may do so in part because their civilian earning potentials are low because of unobservable characteristics. Perhaps offsetting this effect, the services may have been able to observe more about the individuals they

accepted than we can now, selecting the particularly able.⁶ Second, those who leave the military may not be able to step immediately onto the age/earnings profile of their civilian counterparts, and the pattern of wage penalties they suffer may vary with the length of their service. Cooper (1981) offers two pieces of evidence on this issue with regard to retirees: (1) although military retirees ultimately earn more than nonretiree veterans (controlling for hours worked), in the first few years after retirement they earn considerably less; and (2) the later an officer retires the lower his civilian earnings (controlling for grade).⁷

The problem of military officers forming a nonrandom sample of the general population is not a severe one for the DRM. As we noted above (Sec. II), the model's taste parameter will capture, at least in part, permanent differences among officers' civilian earning potentials, and the estimated mean (or location parameter) of the taste distribution will reflect in part any divergence between the average potential earnings of officers and the average in the civilian sector. We say in part because an individual's taste value is assumed to be constant over time in real dollar terms; differences in individual characteristics are generally thought to affect the logarithm of earnings, implying a constant proportional effect over time. In the course of this study we will explore the possibility of altering the DRM formulation by entering the taste parameter as a multiplicative (in earnings) constant, rather than an additive term.

The second problem--that officers may not be able to jump immediately onto the civilian age/earnings profile--is more serious, and is a principal reason for resorting to data on the actual earnings of veterans. Unfortunately, using such data presents its own set of problems.

⁶To control for the possible selectivity of the services, many studies examining the earnings of military retirees have compared them with the earnings of nonretiree veterans. See, for example, Raduchel et al. (1978), Danzon (1980), and Cooper (1981). Goldberg and Warner (1982) compare the earnings of veterans who served military careers of various lengths.

⁷Cooper does not control for self-selection in the timing of retirement, an issue we discuss below.

Among officers at the same point in their military careers, those who choose to leave may do so in part because they perceive their potential civilian earnings to be particularly high, and those who choose to stay may think their potential earnings are low. This is the self-selection problem, which if present would make the average earnings of former officers a biased (upward) estimate of the potential earnings of officers who remained in the military. The difficulty arises because self-selection implies that the later an officer leaves, other things equal, the lower will be on the distribution of potential civilian earnings among all officers who started in the military with him. Thus, the difference between the potential earnings of those who stay and the actual earnings of those who leave may narrow as the year of service increases.*

The self-selection timing problem may limit the usefulness for the current study of data on the actual earnings of veterans. If they were to be used as the basic civilian earnings series, it appears that the model would have to be modified to account directly for self-selection. Another use is possible, however--to determine the size and length of the short-term wage penalty associated with the transition to civilian employment. Information on the wage penalty could be used to modify data on general civilian earnings.

A final issue in the estimation of potential civilian earnings concerns the possibility that an officer who leaves the military may not be able to obtain civilian employment immediately, or to obtain employment in his desired occupation at all. Unemployment plays no explicit role in the DRM, but Gotz and McCall attempt to control for it by multiplying their civilian earnings estimate in each period by one minus the adult unemployment rate, in effect constructing an *expected*

*Goldberg and Warner (1982) attempt to control for individual ability in order to construct a potential civilian earnings series, based on the actual earnings of veterans, that is free of the bias caused by self-selection. Their method requires the assumption, however, that servicemembers know in advance the exact period in which they will leave, an assumption that is explicitly rejected by the DRM formulation. The ACOL-like approach of Goldberg and Warner suggests that the required civilian earnings series could be constructed implicitly in the estimation of an ACOL retention model, but they do not explore this possibility.

civilian earnings series. Similarly, for military pilots Gotz and McCall estimate potential civilian earnings as a weighted average of the general civilian earnings estimate and an estimate of the wages being offered to commercial pilots, the weights being determined by the military pilot cohort size and by airline hiring rates. Both procedures are ad hoc, but with respect to the first it is unfortunately impractical to estimate a separate effect for unemployment. A possible approach to the pilot-hiring problem is the estimation of a separate "taste for flying" distribution, which should be identifiable because the opportunity for flying service has varied with years of service in the military and with the airline hiring rate in the civilian sector.⁹

ESTIMATION

The full details of the estimation procedure are beyond the scope of this paper,¹⁰ and in any event would be of interest only to someone trying to duplicate the estimation and not to a potential user of the estimated model. Some familiarity with the procedure is required, however, for one to understand what gives rise to the computational costs associated with the model, and to appreciate the limitations that practical considerations impose on the model extensions that may be considered. As with the model description in Sec. II, the important elements of the estimation procedure can be described without recourse to mathematical expressions.

The essence of the estimation procedure is the finding of that set of parameter values that makes the model's retention rate predictions most consistent with the observed events. The measure of consistency with the observed events is the likelihood function, which is simply the product of the probabilities that the model attaches to the various events, given a set of parameter values. The value of that function is larger the more the model attaches high probabilities to events that occur frequently and low probabilities to ones that occur infrequently.

⁹Including a taste for flying distribution may also solve an out-of-sample prediction problem encountered by Gotz and McCall, a tendency to underpredict pilots' retirement rates.

¹⁰Gotz and McCall (1984) show the basic mathematical structure of the likelihood function, but give no details of the computational procedures followed.

Thus, the set of parameter values that maximizes the likelihood function provides the best fit to the data. The search for those maximum-likelihood values may either be performed using mathematical optimization techniques or be conducted by the researcher manually through successive trial values. Both methods were used by Gotz and McCall.

The two levels of the model define two steps in the estimation procedure. The first step mirrors the individual-decision level of the model, calculating costs of leaving at various decision points for representative individuals with a broad range of taste values.¹¹ These calculations require values for two parameters, the variance of the random shocks (σ^2) and the discount rate. These are the parameters for which the search for maximum-likelihood values is conducted manually; to begin the process, initial trial values are supplied. Then, for each representative individual we calculate the probability that each observed event could have occurred, given the individual's taste value. How this is done is most easily seen in an example.

Consider a typical officer's decision problem at his first two decision points. If there were no random shocks, he would stay in the first period if his cost of leaving exceeded zero, and leave otherwise. Including random shocks, he will stay if his net return to staying--his cost of leaving *plus* the value of the random shock he receives--exceeds zero, which is to say if the value of the random shock exceeds the negative of his cost of leaving. The *probability* of that occurrence can readily be calculated, given a trial value for the variance of the random shocks and the assumption that the shocks are independent and normally distributed. Similarly, the probability of a leave in the second period is the probability that his second-period shock is less than the negative of his cost of leaving for that period. Of course, he must still be in the service at that point; thus, the probability of the event "stay in the first period, leave in the second" is the *product* of the separate probabilities for the two decisions.

¹¹Recall from Sec. II that the costs of leaving are calculated in a dynamic program.

With the calculation of the event probabilities--one for each event for each of the representative values of the taste parameter--the first step is completed. The search for maximum-likelihood values of the remaining three parameters can take place without recalculating those probabilities, until the process is begun again with new trial values for σ^2 and the discount rate. The remaining three parameters--the location and scale parameters of the taste distribution and the selectivity parameter--together describe the *conditional* distribution of tastes at the first decision point, conditional on the type of commission held by the officer (regular or reserve).

At this point the need for the representative individuals with a broad range of taste values becomes apparent. We are trying to find the distribution of tastes across the population that gives greatest probability to frequent events; thus, we need to know the probabilities of each event for individuals with high, low, and various in-between taste values. Given a set of values for the parameters describing the taste distribution, the *likelihood* of an event is the weighted sum of the probabilities of that event for each of the representative individuals, the weight for each being the density of the taste distribution at the individual's taste value.¹² The product of the likelihoods for all the events is the likelihood function whose maximum is sought. The computational procedure that finds the maximum-likelihood values can be roughly described as an iterative search over alternative values, with additional quantities calculated that aid in narrowing the search quickly. It has the advantages over a manual search that it substitutes computer time for analyst time, and can yield useful measures of the parameter estimates' precision.

The researcher's judgment is substituted for the computer's precision in the search for maximum-likelihood values of the first two parameters (σ^2 and the discount rate) in order to conserve computer time. Each alternative set of values for these two parameters requires a complete recalculation of the costs of leaving and event probabilities for the representative individuals. The mathematical optimization

¹²In the actual procedure, numeric integration replaces the weighted summation that is described.

procedure in effect "tries" a large number of alternative values; if it took over the search for the two parameters that are directly involved in the individual-decision portion of the model, computational costs would increase dramatically. Of course, costs might be reduced if the search for the three taste parameter values were also done manually, but the savings would not be as great because the individual-decision portion would not be affected. At present, computational costs for a single run (one set of the two manual-search parameters) are roughly evenly split between the individual-decision and aggregate portions of the estimation.

The total computational cost is substantial enough that the potential return to any added complication in the model must be weighed carefully against the cost it would impose. Added complication means added parameters; the parameters that might be added fall into three groups. First, and least costly, a parameter could be added to the three already describing the taste distribution. One might think, for example, that the distribution would be more spread and farther to the left for large yeargroups than for small, suggesting the addition of parameters to measure the relationship between yeargroup size and the taste distribution parameters. Second, a parameter might be added that, like the variance of the random shocks, affects individual decisions and is assumed to be the same for all individuals. Suppose, for example, we suspected that the ending of GI Bill eligibility for officers entering after 1976 made those officers less likely to leave the military than their predecessors, and were willing to assume that everyone placed the same monetary value on this change. To estimate that amount we would have to try several alternative values, making the addition of this parameter more expensive than the first, but it would not be necessary to make the many recalculations of event probabilities that a computer search would require. Third, and most costly, another set of distribution parameters might be added. Perhaps it might be thought that individual differences in discount rates are important, requiring the identification of a discount rate distribution. Instead of the 50 (say) representative individuals with their range of taste values, we would then have to calculate event probabilities for roughly $50 \times 50 = 2,500$ individuals, pairing each taste value with each of a range of discount

rates.¹³ We do not contemplate adding a discount rate distribution, but have considered incorporating a distribution of "tastes for flying," which as noted above might solve a problem that Gotz and McCall encountered in predicting retention rates for pilots in the retirement-eligible years.

PREDICTION

Prediction with the DRM is not the straightforward reversal of the estimation process that it is with most regression models, differing in four important respects. First, every set of predictions requires a model simulation. There are no simple answers to such seemingly simple questions as what would be the effect of a pay increase because the model's parameters describe officers' preferences, not their average response to specific stimuli. Second, the model's predictions are explicitly dynamic. Its answer to the simple pay-increase question will be not one, but several; one for each year in the future. Third, because the model's retention rates are backward looking, data on historical compensation and personnel policies are as important a part of the input for projections as are data on future policies. Fourth, the rich picture of the officer's military career that is part of the model makes its predictions too detailed for easy analysis without the aid of an inventory projection model or, at the very least, some input of the approximate numbers of officers in various states. These differences make prediction with the DRM a more complex process than it is with the ACOL model, for example, but as shown in Sec. III they carry the offsetting reward of much greater model flexibility. Note in addition that the complexity will for the most part be transparent to the DRM's user, except as it results in longer computation times than would be common with a simpler model.

In discussing estimation above, we traced out the steps necessary to produce predictions of event probabilities, the underlying elements of retention rate predictions. With the model's parameter values known, the iterative searches in the estimation process are no longer required,

¹³The actual number would be somewhat less because techniques exist for numerically evaluating a double integral that do not require as dense a "grid" of values as is needed in the one-dimension case. The order of magnitude of the increase, however, would be as indicated.

but the probability of a stay in each period along every feasible path must be calculated,¹⁴ including periods many years beyond the initial decision point. (These calculations will be performed by a computer, of course, not by the DRM's user.) With the ACOL model, in contrast, once the ACOL for a particular YOS point has been calculated, the retention rate is readily produced from the regression equation. The effect of a pay increase can even be calculated analytically from the ACOL slope parameter, so long as the increase does not alter the horizon (the planned leave point) of the marginal officer. The DRM, however, recognizes that the effects of the pay increase on every officer, whether on the margin or not, are important in determining aggregate retention rates. Thus, the DRM requires a simulation of the pay-increase effect for officers with high, low, and in-between taste values, and the aggregation across them using the estimated taste distribution.

The dynamic nature of the DRM's predictions make them more complex than those of regression models such as the ACOL, but this is simply a reflection of the complexity of reality. A pay increase this year should be expected to have a different effect on retention at the ten-year point next year than it would at the same point five years hence. The groups of officers at the two points will be different because they will have faced different policies up to those points, not the least of those differences being the greater pay that the later group will have received in the five interim years. The DRM can produce steady-state predictions, of course, but in a constantly changing world steady states do not appear very relevant. It is interesting to note that although the ACOL's predictions appear to be steady-state they are not, reflecting instead the partial adjustments of the single snapshot from which the model's parameters were estimated.

That projecting retention rates requires historical data on the policies and opportunities facing officers is a natural consequence of the DRM's backward-looking retention rates, the reasons for which should not require further discussion at this point. The need creates a

¹⁴As a practical matter, paths that occur with very low probability, such as ones leading to promotion to O-6 at 13 years of service, can be ignored.

problem, however, that deserves note. Given the availability of reliable data only since 1973, the historical record for many yeargroups is not complete. Making predictions for the group of officers that will reach retirement eligibility in 1985, for example, requires data on the environments they faced every year since their first decision points, which may have been as early as 1968. Gotz and McCall filled in the missing information by assuming that it matched the conditions facing the earliest yeargroup for which they had complete information. We are exploring the possibility of improving on that assumption by using available information on the proportion of each yeargroup still in the service when the group reached the beginning of our observation period. Should this attempt not prove fruitful, our longer historical period than was used by Gotz and McCall will at least ensure that the problem of missing information will not affect the majority of relevant yeargroups. Moreover, with each passing year the problem will become less important.¹⁵

The fourth difference between the DRM and simpler regression models is illustrated by considering the problem of predicting a retention rate for officers in their eleventh year of service in 1985. Some of these officers will have been promoted to O-4 a year or more earlier, some will just have been promoted in due course, and some will have been passed over for due-course promotion. Presumably, the retention rates for these three groups will differ, as they face different future opportunities and the third group is in a different grade than the other two. The DRM will provide a separate prediction for each group, which would prove too cumbersome for easy analysis. Gotz and McCall used rough estimates of the proportions of the yeargroup that would fall in each promotion group, using those estimates to form the weighted sum of the three retention rates. A similar procedure could be used to aggregate across officer communities. Much simpler in both cases, however, would be to derive the numbers in each group from an inventory

¹⁵As time passes, the yeargroups of interest will tend more and more to have made their first retention decisions in the period of the historical data assembled for estimation. In addition, each new year's data on promotion, pay, etc., can be added to the historical data, helping to maintain the model's reliability even in the absence of reestimation. The computer programs developed in this study will be adapted for use by DMDC personnel in generating the new data.

projection model (IPM) integrated with the DRM, perhaps bypassing altogether the computation of average retention rates and basing whatever analysis was required on the officer inventories themselves.

An IPM integrated with the DRM would offer another advantage in making predictions. One of the inputs the DRM requires is the set of transition probabilities among states, a key element of which is the set of promotion probabilities. For the historical period, these promotion probabilities can be derived from the data, based on actual promotion rates (as the individual officer would likely calculate his chances of promotion). For the projection period, however, we would want to base the probabilities on announced or planned service policies with respect to promotion opportunity and/or timing, and might wish to examine the effects of changes in either of those policy parameters. Translating the policy parameters into promotion probabilities requires information on inventories, making it a natural role for an IPM.

A typical projection run of an integrated DRM-IPM, then, might look something like this. Suppose we were examining the effects of a pay increase, and knew the service's plans with respect to promotion opportunities and timing. The IPM would be designed to accept these as input parameters, and would derive the associated promotion probabilities, feeding them to the DRM. The analyst would alter the military pay inputs to the DRM, which would then calculate a detailed set of retention rates. The IPM would derive the resultant inventories, and perhaps calculate aggregate retention rates by year of service. If the inventories exceeded grade table limits at some future point, the analyst might want to adjust promotion timings (or the IPM could be designed to derive timings given opportunities and grade table limits). If the changes in timings were fairly small, the process could stop at this point: if they were large, another pass through the DRM would be desirable to determine the retention effects to the timing changes.

The DRM cannot reside in a programmable pocket calculator, but this does not mean that it cannot be made user-friendly. Making it so requires only the writing of some computer code, and decisions as to which potential policy changes should be readily accommodated through simple changes in the model inputs, and which are so unlikely that simulating their effects could best be left to someone capable of modifying the computer code.

V. CONCLUSION

The Dynamic Retention Model is a versatile tool for the analysis of policy actions affecting military officers' retention. It models the decision process over time of officers facing an uncertain environment, rather than describing average responses to specific stimuli, which allows it to predict plausibly the effects of policy changes that have no historical analogues. By explicitly recognizing that individuals differ persistently in their satisfaction with military service, the DRM can trace the effects of past policies on future retention rates.

Although the DRM is in reality a fairly simple depiction of the determinants of military retention, its seeming complexity has apparently stood in the way of its general acceptance. This paper has attempted to dispel the misperception of the DRM as too complex and difficult to understand for everyday use. As we have shown, the heart of the model does not lie in complex mathematical expressions, but rather in an intuitively appealing common-sense picture of officers' rational behavior and a straightforward tracing out of the aggregate implications of that behavior. The model rests on a number of assumptions, which we have made explicit, and if the number seems large it is only because no comparable listing exists for competing models. By building the DRM step by step through its assumptions, we have demonstrated that the more commonly used Annualized Cost of Leaving model is a limiting special case embodying an important logical inconsistency.

Two other criticisms that have been raised against the DRM are that it requires a great deal of data and is difficult to estimate. With the experience provided by the developmental work of Gotz and McCall, however, these problems are not severe. The data requirements are extensive but fairly straightforward, principally involving the generation and use of longitudinal master files for officers. The Defense Manpower Data Center has recently constructed a number of such files for various studies, including files used by the Fifth Quadrennial Review of Military Compensation. DMDC will be provided with computer

programs for updating the model's inputs as additional years of data become available. Estimation is more complex than for a model such as the ACOL, but the DRM is specifically designed to be estimated by maximum-likelihood, making the estimation procedure quite easy to understand. Workable computer code for estimation has been developed by Gotz and McCall, allowing the current project to work on improving efficiency and adding refinements.

Given the feasibility of estimating and using the Dynamic Retention Model, the important question becomes whether it really is capable of dealing with a substantially greater variety of policy changes than a model such as the ACOL. The examples given in Sec. III demonstrate that the DRM is, indeed, able to yield plausible predictions in cases where the ACOL is not. Changes in the structure of compensation as well as its level, in important personnel policy parameters, and in the structure of retirement benefits can readily be accommodated by the DRM. The parameters of the ACOL, in contrast, are strongly determined by the structures of compensation, retirement, and personnel policies in effect during the period for which they are estimated. The insights provided by the DRM allow us to show that the ACOL will incorrectly predict the effects of even very minor changes in policy, although not to determine how important the errors will be in practice.

No model is a perfect depiction of reality, of course. In the process of describing the DRM we have pointed out areas in which changes could be made to improve the model's predictive accuracy. Several of these center on the estimation and treatment of civilian earnings opportunities: incorporating data on the actual earnings of officers who leave while controlling for self-selection, determining the extent to which veterans' earnings reflect a discounting of military experience on the part of civilian employers, and treating more explicitly unmeasured differences in individual earnings opportunities. Another area concerns the treatment of additional service obligations incurred during the period of the initial active duty service commitment; this is less a possible improvement than it is a necessary response to the unavailability of obligation information for at least one service. A third potential improvement is the decomposition of individuals' tastes for military service into two components, one reflecting satisfaction

with a particular type of duty (such as flying), the other embodying the general quality of the match between the individual and the characteristics of military versus civilian employment. We will discuss these issues more fully in a forthcoming paper.

Finally, we should point out that the complexity of the DRM is not something that need concern the user of the models being developed in this study. Once the DRM's parameters have been estimated, they will be imbedded in a projection version of the model that will allow for easy modification of its data inputs, including military and civilian pay levels. Programs provided to DMDC will generate additional historical data with each passing year. A user-friendly IPM will accept scenarios specifying service plans (or possible alternatives) with respect to promotion and tenure policies, and provide the primary user interface with the DRM. Reestimation of the model's parameters will require more technical skills than using the model, but might not be necessary for many years because the estimated parameters are not strongly dependent on existing compensation levels and personnel practices. When necessary, reestimation will be facilitated by programs developed in this study. Day-to-day operation of the projection model will not require any extensive training, and will be made easy by a user's manual.

BIBLIOGRAPHY

- Brown, Charles, "Equalizing Differences in the Labor Market," *Quarterly Journal of Economics*, Vol. 94, No. 1, February 1980, pp. 112-134.
- Chipman, M. D., and R. H. Mumm, *Forecasting Naval Enlisted Retention Behavior under Alternative Retirement Systems*, Naval Personnel Research and Development Center, Tech. Rep. 79-4, San Diego, CA, November 1978.
- Cooper, Richard V. L., *Military Retirees' Post-Service Earnings and Employment*, The Rand Corporation, R-2493-MRAL, February 1981.
- Danzon, Patricia Munch, *Civilian Earnings of Military Retirees*, The Rand Corporation, R-2353-MRAL, March 1980.
- Goldberg, Matthew S., and John T. Warner, *Post-Service Earnings of Military Veterans*, Center for Naval Analyses, CRC 472, Arlington, VA, October 1982.
- Gotz, Glenn A., and John J. McCall, *A Sequential Analysis of the Air Force Officer's Retirement Decision*, The Rand Corporation, N-1013-1-AF, October 1979.
- , *Estimating Military Personnel Retention Rates: Theory and Statistical Method*, The Rand Corporation, R-2541-AF, June 1980.
- , "Estimation in Sequential Decisionmaking Models: A Methodological Note," *Economics Letters*, Vol. 6, 1980, pp. 131-136.
- , "Sequential Analysis of the Stay/Leave Decision: U.S. Air Force Officers," *Management Science*, Vol. 29, No. 3, March 1983, pp. 335-351.
- , *A Dynamic Retention Model for Air Force Officers*, The Rand Corporation, R-3028-AF, December 1984.
- Johnson, George E., "Intermetropolitan Wage Differentials in the United States," in Triplett (1983).
- Leibowitz, Arleen, "Fringe Benefits in Employee Compensation," in Triplett (1983)
- Raduchel, William J., et al., "Post Retirement Income and Earnings of Military Personnel Who Retired from 1970 to 1974," in *Supplementary Papers of the President's Commission on Military Compensation*, U.S. Government Printing Office, Washington, D.C., 1978.

Smeeding, Timothy M., "The Size Distribution of Wage and Nonwage Compensation: Employer Cost versus Employee Value," in Triplett (1983).

Smith, Robert S., and Ronald G. Ehrenberg, "Estimating Wage-Fringe Trade-Offs: Some Data Problems," in Triplett (1983).

Triplett, Jack E., editor, *The Measurement of Labor Cost*, The University of Chicago Press, New York, 1983.

Warner, John T., *Alternative Military Retirement Systems: Their Effects on Enlisted Retention*, Center for Naval Analyses, CRC 376, Arlington, VA, September 1979 (unpublished mimeo dated 18 December 1978).

-----, *Military Compensation and Retention: An Analysis of Alternative Models and a Simulation of a New Retention Model*, Center for Naval Analyses, CRC 436, Arlington, VA, August 1981.

-----, and Matthew S. Goldberg, "The Influence of Non-Pecuniary Factors on Labor Supply: The Case of Navy Enlisted Personnel," *Review of Economics and Statistics*, Vol. 64, No. 1, February 1984, pp. 26-35.

END

FILMED

8-85

DTIC